

RECEIVED
CENTRAL FAX CENTER

OCT 10 2006

REMARKSI. Introduction

In response to the Office Action dated July 10, 2006, claims 1, 11 and 21 have been amended. Claims 1, 3-10, 13-21 and 23-30 remain in the application. Re-examination and re-consideration of the application, as amended, is requested.

II. Claim Amendments

Applicant's attorney has made amendments to the claims as indicated above. These amendments were made solely for the purpose of clarifying the language of the claims, and were not required for patentability or to distinguish the claims over the prior art.

III. Double Patenting Rejection

On page (2) of the Office Action, claims 1 and 3-30 were provisionally rejected on the ground of nonstatutory double patenting over claim 1-33 of copending Application No. 10/807,871.

Applicant's attorney notes the provisional nature of these rejections, and will respond substantively to the rejections once allowable claims have been identified.

IV. Statutory Subject Matter Rejection

On page (3) of the Office Action, claims 1, 3-11, 13-21 and 23-30 were rejected under 35 U.S.C. §101 because the claimed invention is allegedly directed to non-statutory subject matter.

Applicant's attorney has amended claims 1, 11 and 21 as indicated above to overcome the rejections, but nonetheless traverses the rejections.

Contrary to the assertion of the Office Action, Applicant's attorney submits that the generation of cardinality estimates that are then used to determine an optimal query execution plan is, in fact, limited to a substantial practical application, namely database management, so as to optimize queries performed by a database management system. Indeed, the Applicant's claimed invention recites a specific instance, not just a concept or idea, which indicates that the final result achieved by the claimed invention is useful, tangible and concrete. Consequently, Applicant's attorney requests withdrawal of the rejection.

However, should issues still remain in this regard, Applicant's attorney requests that the Examiner indicate how the rejection can be overcome, in accordance with the directives of the Examination Guidelines for Computer-Related Inventions. See Guidelines II M.P.E.P. § 2106.

Specifically, should it be necessary, the Applicant's attorney requests that the Examiner identify features of the invention that would render the claimed subject matter statutory if recited in the claim. See Guidelines IV, M.P.E.P. § 2106.

V. Prior Art Rejections

A. The Office Action Rejections

In paragraphs (1)-(2) of the Office Action, claims 1, 3-5, 11, 13-15, 21 and 23-25 were rejected under 35 U.S.C. §102(e) as being anticipated by U.S. Patent No. 6,026,391 (Osborn). In paragraphs (3)-(4) of the Office Action, claims 6-10, 16-20 and 26-30 were rejected under 35 U.S.C. §103(a) as being obvious in view of the combination of Osborn in view of U.S. Patent No. 6,496,819 (Bello).

Applicant's attorney respectfully traverses these rejections.

B. The Applicant's Claimed Invention

Applicant's claimed invention, as recited in independent claims 1, 11 and 21, is generally directed to a method of optimizing execution of a query that accesses data stored on a data store connected to a computer. Claim 1 is representative and recites the steps of generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query, using the generated cardinality estimates to determine an optimal query execution plan for the query, and executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.

C. The Osborn Reference

Osborn describes systems and methods for estimating query response times in a computer system. A query performance prediction ("QPP") module is provided as part of an applications layer residing on respective user stations operating in conjunction with a centralized host computer system. The QPP module correlates estimated system cost information for a database query provided from a database management system associated with the computer system with statistics compiled from previous queries in order to estimate the system response time. In particular, an estimated CPU time for executing the present query is derived by extrapolating actual CPU times recorded for past queries having the closest estimated costs for accessing the same, or similar, tables

and items in the database, using a form of a “nearest neighbor” algorithm to match the present query to either identical or statistically closest past queries. The estimated CPU time for the present query is then multiplied by a current ratio of total elapsed times-to-CPU times for the system to produce an estimated total elapsed time for responding to the query.

D. The Bello Reference

Bello describes a method and system for processing queries. Specifically, techniques are provided for handling a query that does not reference a particular materialized view, where the query requires access to values from a particular column not contained in the materialized view. A technique is also provided for processing a query that does not reference a particular materialized aggregate view, where the materialized aggregate view specifies an outer join between a child table and a parent table and the query specifies a particular type of join between the child table and the parent table, where the particular type of join is one of an inner join, an anti-join and a semi-join. The query is rewritten to produce a rewritten query that accesses the materialized aggregate view to produce data required by the query. A technique is also provided for processing a query that does not reference a particular materialized view and that specifies that results are to be grouped by a first set of one or more columns, where the materialized view reflects data that is grouped by a second set of one or more columns. If each column in the first set of columns either exactly matches a column in the second set of columns, is functionally dependent on another column in the first set of columns, or is functionally dependent on a column in the second set of columns, then the query is rewritten to produce a rewritten query that references the materialized view.

E. Applicant's Claimed Invention Is Patentable Over The Cited References

Applicant's claimed invention is patentable over the cited references, because it includes a combination of limitations not taught or suggested by the cited references, taken individually or in combination. Specifically, neither reference teaches or suggests the steps or elements of the independent claims comprising “generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query,” and “using the cardinality estimates to determine an optimal query execution plan for the query.”

Nonetheless, Osborn is cited by the Office Action as teaching all of the steps or elements of the independent claims 1, 11 and 21. The portions of Osborn cited by the Office Action are set forth below:

Osborn: Col. 6, lines 24-64

In particular, the cost optimizer 42 generates a set of potential execution plans (not shown) for executing the respective query SQL statement (40) based on the available access paths, and estimates the relative "cost" of each potential execution plan based on the data distribution and storage characteristics for the respective tables, clusters and indexes to be used. Execution plans with greater relative costs generally take more time to execute than those with smaller relative costs, although the relationship between cost and execution time is by no means linear, and depends on many other factors. An exemplary preferred cost optimizer is provided in the Oracle7.3.TM. version of the Oracle.RTM. database product produced and distributed by Oracle Corporation, Redwood City, Calif., USA.

The cost optimizer 42 compares the estimated costs of the potential execution plans and returns the smallest estimated cost 44, along with a result set 45 representing the selected execution plan for the input query 40, to a query performance prediction ("QPP") module 46 residing within the database applications layer 36 of the respective user station 24. The QPP module 46 compares the estimated cost 44 and result set 45 for the present query 40 to the recorded estimated costs and result sets of past queries, searching out those past queries having the same or similar estimated costs for accessing the same, or similar, tables and items in the database located in a query statistics cache 48 associated with the database applications layer 36.

More particularly, referring additionally to FIG. 3, each user station 24 records pertinent information for each new user query, which are collected and maintained by the host computer 22 in a query history 50. In a presently preferred embodiment, the query history 50 includes: the query ID 58, the date 60 and time 62 at which the query was executed; the user ID 64 of the user station 24 submitting the query; an identification of the result set 66 for the query--i.e., the particular table(s) and column(s) that were accessed; whether the query was satisfied from a pre-computed summary table 68; the estimated cost 70 for the query; the actual CPU time 72 used in running the query; and the estimated time 73 and actual total elapsed time 74 required for the system 20 to respond to the query.

Osborn: Col. 7, lines 6-35

Referring additionally to FIG. 4, the QPP module 46 compares 80 the cost estimate 44 and result set 45 for the present query 40 to the recorded estimated costs and result sets of past queries stored in the query statistics cache 48. If an exact match is found between the present query 40 and a recorded past query 58 stored in the cache 48 (i.e., wherein both queries have the same estimated costs for the same result sets), the QPP module 46 selects the recorded actual CPU time 72 of the matching past query 58 as an estimated CPU time 82 for the present query 40.

If no exact match is found, a nearest neighbor algorithm 83 is employed to extrapolate an estimated CPU time 84 based on a weighted average of CPU times for the closest matching stored queries--i.e., from the "nearest neighbors" of the present query 40 based on the respective recorded estimated costs and result sets of the past queries. Notably, if a past query having an identical result set was executed using a summary table, the estimated response time can be approximated to a selected (e.g., minimum) value. A preferred methodology for the use of summary tables is provided in U.S. patent application Ser. No. 08/962,533, entitled "Summary Table Query Routing," filed on the same day as the present application and fully incorporated herein by reference. A preferred methodology for creating summary tables is provided in U.S. patent application Ser. No. 08/962,029, entitled "Summary Table Management In A Computer System," which was also filed on the same day as the present application and which is also fully incorporated herein by reference.

Nothing in the above portions of Osborn can fairly be said to represent "generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query," and "using the cardinality estimates to determine an optimal query execution plan for the query."

Instead, the above portions of Osborn merely state that the relative cost of each potential execution plan is estimated based on the data distribution and storage characteristics for the respective tables, clusters and indexes to be used. Moreover, the query performance prediction ("QPP") module of Osborn compares the estimated cost and result set for the query to a query history for past queries, searching out those past queries having the same or similar estimated costs for accessing the same, or similar, tables and items in the database located in a query statistics cache. This query history not only includes the estimated cost for the query, but also identifies whether the query was satisfied from a pre-computed summary table. If a past query having an identical result set was executed using a summary table, the estimated response time can be approximated to a selected (e.g., minimum) value.

However, nowhere does Osborn describe the use of cardinality estimates generated using statistics of automatic summary tables that vertically overlap the query. As noted in Applicant's specification, and as acknowledged by the Office Action, automatic summary tables are pre-computed queries or materialized views. Also as noted in Applicant's specification, although not acknowledged by the Office Action, an automatic summary table vertically overlaps a query when the set of predicates applied by the automatic summary table is a subset of the predicates required by the query.

However, there is no discussion of vertically overlapping automatic summary tables in either Osborn. Moreover, the Office Action ignores this limitation when asserting that Osborn teaches all the limitations of Applicant's claims, which is improper.

Bello fails to overcome the deficiencies of Osborn. Recall that Bello was cited only against dependent claims 6-10, 16-20 and 26-30, and only for teaching the application of predicates to automatic summary tables.

Consequently, the Osborn reference alone, and the Osborn reference combined with the Bell reference, does not teach or suggest Applicant's invention. Moreover, the various elements of Applicant's claimed invention together provide operational advantages over the cited references. In addition, Applicant's invention solves problems not recognized by the cited references.

Thus, Applicant submits that independent claims 1, 11 and 21 are allowable over Osborn and Bello. Further, dependent claims 3-10, 13-20 and 23-30 are submitted to be allowable over Osborn and Bello in the same manner, because they are dependent on independent claims 1, 11 and 21, respectively, and because they contain all the limitations of the independent claims.

VI. Conclusion

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicant's undersigned attorney.


Respectfully submitted,

GATES & COOPER LLP
Attorneys for Applicant

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: October 10, 2006

GHG/cf

By: 
Name: George H. Gates
Reg. No.: 33,500

G&C 30571.272-US-U1